Intel® Omni-Path Fabric Software

Release Notes for 10.0.1

April 2016

Contents

Tables

§

1.0 Overview of the Release

1.1 Introduction

This document provides a brief overview of the changes introduced into the Intel® Omni-Path Software by this release. References to more detailed information are provided where necessary. The information contained in this document is intended as supplemental information only; it should be used in conjunction with the documentation provided for each component.

These Release Notes list the features supported in this software release, open issues, and issues that were resolved during release development.

1.2 Audience

The information provided in this document is intended for installers, software support engineers, service personnel, and system administrators.

1.3 Software License Agreement

This software is provided under license agreements and may contain third-party software under separate third-party licensing. Please refer to the license files provided with the software for specific details.

1.4 If You Need Help

Technical support for Intel® Omni-Path products is available 24 hours a day, 365 days a year. Please contact Intel Customer Support or visit www.intel.com for additional detail.

1.5 New Features in this Release

This release contains no new features, it is a maintenance release that contains bug fixes and performance improvements.

1.6 Supported Features

- The list of supported operating systems is in Table 1-2.
- The list of supported hardware is in Table 1-5.
- Coexistence with Intel® True Scale Architecture. This release supports True Scale hardware serving as an Infiniband* storage network with the Intel® Omni-Path hardware used for computing. Note that connecting a True Scale adapter card to an Omni-Path switch, or vice-versa, is not supported. For more details on this feature, refer to *Intel® Omni-Path Fabric Host Software User Guide.*
- Supports Dual Rail: Two Intel® Omni-Path Host Fabric Interface (HFI) cards in the same server connected to the same fabric.

- Supports Dual Plane: Two HFI cards in the same server connected to separate fabrics.
- Limited validation testing performed on network storage file systems:
    — NFS over TCP/IP
- Active Optical Cables
- Intel® Omni-Path Boot over Fabric
- MPI applications are provided in a stand-alone rpm package.
- Intel® Xeon® v4 processor (codename Broadwell) support
- KNL support (B0 technology preview)
- Monitored Intel® Omni-Path Host Fabric Interface
- DHCP and LDAP supported on Intel® Omni-Path Edge Switch 100 Series and Intel® Omni-Path Director Class Switch 100 Series hardware.

1.7 Release Packages

There are two Intel® Omni-Path Fabric Software packages:

- Basic for compute nodes
- IFS for the management node

The Basic package includes:

- Software that installs the following packages to the distribution OpenFabrics Alliance* (OFA):
    — libibumad is based on the RHEL* or SLES* release package. It contains Intel patches that support Intel® Omni-Path Architecture (Intel® OPA) technology.
    — ibacm is the latest upstream code applied with RHEL* patches.
    — hfi1-firmware, hfi1-psm, hfi1-diagtools-sw, libhfi1verbs
    — Open MPI built for verbs and PSM2 using gcc, Intel, and PGI compilers.
    — MVAPICH2 built for verbs and PSM2 using gcc, Intel, and PGI compilers.
    — mpitests
    — mpi-selector
    — GASnet
    — openSHMEM
    — srptools (includes the latest upstream code)
    — Firmware files listed in Table 1-1.
- compat-rdma which delivers kernel changes based on the OFA version. The components installed are the hfi1 driver and Intel-enhanced versions of other kernel packages. See the *Intel® Omni-Path Fabric Software Installation Guide* for details.

    Note: In the Intel® Omni-Path Software package for RHEL* 7.2, the hfi1 driver and ifs-kernel-updates are supplied as a smaller package.

The IFS package includes the Basic package plus:

- Fabric Manager, which allows comprehensive control of administrative functions using a mature Subnet Manager. Fabric Manager simplifies subnet, fabric, and individual component management, easing the deployment and optimization of large fabrics.
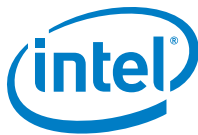
- Fabric Suite FastFabric Toolset, which enables rapid, error-free installation and configuration of Intel® OPA host software and management software tools, as well as simplified installation, configuration, validation, and optimization of HPC fabrics. For details, refer to the Fabric Suite FastFabric documentation in Table 1-6.

1.8 Firmware Files

This release of the Intel® Omni-Path Software contains the firmware files listed in Table 1-1.

Table 1-1. Firmware Files

Description	File Name	Version
Intel® Omni-Path Boot over Fabric	hfi1_fabric_boot_v1.8.rom	v1.8
HFI1 UEFI Option ROM	HfiPcieGen3_0x15.rom	0x15
HFI1 SMBus Microcontroller Firmware (Thermal Monitor)	hfi1_smbus.fw	10.0.0.0.696
HFI1 8051 Firmware	hfi1_dc8051.fw	v0.36

1.9 Operating Systems

This release of the Intel® Omni-Path Software supports the operating systems listed in Table 1-2.

Table 1-2. Operating Systems Supported

Operating System	Update/ SP	Kernel Version
Red Hat* Enterprise Linux* (RHEL*) 7.1 X86_64	Update 1	3.10.0-229.el7.x86_64
CentOS* 7.1 X86_64	N/A	3.10.0-229.el7.x86_64
Scientific Linux* 7.1 X86_64	N/A	3.10.0-229.el7.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.2 X86_64	Update 2	3.10.0-327.el7.x86_64
SUSE* Linux* Enterprise Server (SLES*) 12 X86_64	N/A	3.12.28-4-default
SUSE* Linux* Enterprise Server (SLES*) 12.1 X86_64	Service Pack 1	3.12.49-11.1-default

1.10 Parallel File System

The following parallel file system has been tested with this release of the Intel® Omni-Path Software:

- IBM* General Parallel File System (GPFS) version 4.0.1
  — RHEL* 7.1 and 7.2

Refer to the *Intel® Omni-Path Fabric Performance Tuning User Guide* for details on optimizing parallel file system performance with Intel® Omni-Path Software.
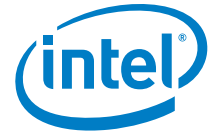
1.11 Compilers

1.11.1 MPI

This release supports the following MPI implementations:

Table 1-3. MPI Compilers

MPI Implementation	Runs Over	Compiled With
Open MPI 1.10.0	Verbs	GCC
	PSM2	GCC, Intel
MVAPICH2-2.1	Verbs	GCC
	PSM2	GCC, Intel
Intel® MPI 5.1.3	Verbs	GCC
	PSM2	GCC, Intel

1.11.2 MVAPICH2 and Open MPI

MVAPICH2 and Open MPI have been compiled for PSM2 to support the following versions of the compilers:

Table 1-4. MVAPICH2 and Open MPI

Compiler	Linux* Distribution	Compiler Version
(GNU) gcc	RHEL* 7	gcc (GCC) 4.8.2 20140120
(GNU) gcc	RHEL* 7.1	gcc (GCC) 4.8.3 20140911
(GNU) gcc	RHEL* 7.2	gcc (GCC) 4.8.5 20150623 (Red Hat* 4.8.5-4)
(GNU) gcc	SLES* 12	gcc (SUSE* Linux*) 4.8.3 20140627
(GNU) gcc	SLES* 12 SP 1	gcc (SUSE* Linux*) version 4.8.5
(Intel) icc	RHEL* 7	icc (ICC) 15.0.1
(Intel) icc	RHEL* 7.1	icc (ICC) 15.0.1
(Intel) icc	RHEL* 7.2	icc (ICC) 15.0.1
(Intel) icc	SLES* 12	icc (ICC) 15.0.1
(Intel) icc	SLES* 12 SP 1	icc (ICC) 15.0.1

See ID 132345 for expected behavior that may be observed when using Open MPI with the SLURM PMI launcher and PSM2.

1.12　Hardware

Table 1-5 lists the hardware supported in this release.

Table 1-5.　Hardware Supported

Hardware	Description
Intel® Xeon® v3 Processor	Haswell CPU-based servers
Intel® Xeon® v4 Processor	Broadwell CPU-based servers
Intel® Omni-Path Host Fabric Interface 100HFA016 (x16)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Host Fabric Interface 100HFA018 (x8)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Switch 100SWE48Q	Managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48U	Externally-managed 48-port Edge Switch
Intel® Omni-Path Director Class Switch 100SWD24	Director Class Switch 100 Series, up to 768 ports
Intel® Omni-Path Director Class Switch 100SWD06	Director Class Switch 100 Series, up to 192 ports

1.13　Installation Requirements

1.13.1　Software and Firmware Requirements

Table 1-2 lists the operating systems supported by this release. Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for the required packages.

1.13.2　Installation Instructions

There are two Intel® Omni-Path Fabric Software packages:

- `IntelOPA-IFS.<distro>-x86_64.<version>.tgz` for the management node.
- `IntelOPA-Basic.<distro>-x86_64.<version>.tgz` for compute nodes.

The packages in the tgz file are RPMs. Installing individual RPMs is not supported in the 10.0.1 release.

Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for related software requirements and complete installation procedures. Refer to the *Intel® Omni-Path Fabric Hardware Installation Guide* for related firmware requirements.

1.13.2.1　Disable SRP Daemon Autostart

By default, RHEL* 7.1 is configured to autostart the SRP daemon on every node. This daemon performs a fabric search for SRP target devices every minute, which may impact the performance of HPC applications, especially in larger fabrics.

Intel recommends you disable the SRP daemon on all compute nodes and any service nodes. Typically, the SRP daemon only needs to be enabled on filesystem server nodes that are directly accessing block storage devices.

When installing either the IntelOPA-Basic or IntelOPA-IFS packages, there is a prompt: `Enable SRP initiator autoload?` Intel recommends you use the default option `no` on all compute notes and any service nodes.
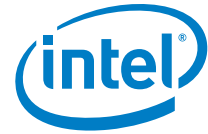
When installing using other means (such as provisioning systems or direct installation of the rpms), you must manually disable the SRP daemon autostart. After installing the `rdma-*-noarch.rpm` package, edit the `/etc/rdma/rdma.conf` file. The default setting is: `SRP_LOAD=yes`. Intel recommends you change this to `no`.

1.14 Product Constraints

This release has the following product constraints:

- This release includes HFI UEFI version 0x15.
  Intel does not recommend downgrading to version 0x14 bcause it may cause a hang condition with some servers on boot. Contact Intel Customer Support for more information.

1.15 Product Limitations

This release has the following product limitations:

- The embedded version of the Fabric Manager supports a maximum of 100 nodes within a fabric. This is due to the limited memory and processing resources available in the embedded environment.

- The UEFI Option ROM and Intel® Omni-Path Boot over Fabric are not signed, and therefore do not support secure boot. In-band PXE boot over Omni-Path (unsecured) is available.

- Short Term History should only be enabled among FMs that are running the same release of software. It should **not** be enabled with FMs running a combination of software releases, for example, release 10.0.0 and pre-production 10.0.0 versions. Prior to starting FMs, short term history files from previous or subsequent software releases should be deleted from the short term history directory:
  `<ShortTermHistory>/<StorageLocation>`
  FMs running short term history with release 10.0 should also have CPUs that are little-endian.
  Short Term History is enabled via
  `configuration: <ShortTermHistory>/<Enable> == 1`

- PA Failover should only be enabled among FMs that are running the same release of software. It should **not** be enabled with FMs running a combination of software releases, for example, release 10.0.0 and pre-production 10.0.0 versions.
  PA Failover is enabled via
  `configuration:<PM>/<ImageUpdateInterval> > 0`

1.16 Documentation

Table 1-6 lists the end user documentation for the current release.

Documents are available at the following URLs:

- Intel® Omni-Path Switches Installation, User, and Reference Guides
  http://www.intel.com/SUPPORT/NETWORK/OMNI-SWTCH100/SB/CS-035856.HTM

- Intel® Omni-Path Host Fabric Interface Installation, User, and Reference Guides (includes software documents)
  http://www.intel.com/support/network/omni-adptr/sb/CS-035857.htm

- Drivers and Software (including Release Notes)
  https://downloadcenter.intel.com/

Table 1-6.   Related Documentation for this Release

Document Title
Hardware Documents
Intel® Omni-Path Fabric Switches Hardware Installation Guide
Intel® Omni-Path Fabric Switches GUI User Guide
Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide
Intel® Omni-Path Edge Switch Platform Configuration Reference Guide
Intel® Omni-Path Fabric Managed Switches Release Notes
Intel® Omni-Path Fabric Externally-Managed Switches Release Notes
Intel® Omni-Path Host Fabric Interface Installation Guide
Fabric Software Documents
Intel® Omni-Path Fabric Software Installation Guide
Intel® Omni-Path Fabric Suite Fabric Manager User Guide
Intel® Omni-Path Fabric Suite FastFabric User Guide
Intel® Omni-Path Fabric Host Software User Guide
Intel® Omni-Path Fabric Suite Fabric Manager GUI Online Help
Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide
Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide
Intel® Omni-Path Fabric Performance Tuning User Guide
Intel® Omni-Path Host Fabric Interface Platform Configuration Reference Guide
Intel® Omni-Path Fabric Software Release Notes
Intel® Omni-Path Fabric Manager GUI Release Notes
Intel® Omni-Path Storage Router Design Guide
Building Lustre* Servers with Intel® Omni-Path Architecture Application Note
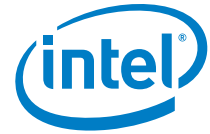
2.0 Issues

2.1 Introduction

This section provides a list of the resolved and open issues in the Intel® Omni-Path Software.

2.2 Resolved Issues

Table 2-1 lists issues that are resolved in this release.

Table 2-1. Resolved Issues (Sheet 1 of 2)

ID	Component	Description	Resolved in Release
130454	HFI Host Driver	Unable to enable pcie error reporting: -5	10.0.1
131176	Host Fabric Software/ MPI	Intel MPI is missing shortcut switch to enable PSM2.	10.0.1
131646	Host Fabric Software	SRP module works fine on connection setup and data transfer, but hangs on module removal 'modprobe -r ib_srp'.	10.0.1
131651	HFI Host Driver	The UEFI driver cannot be used in combination with BIOS Refresh release for Intel® Xeon® Processor E5 v4 (codename Broadwell).	10.0.1
131918	HFI Host Fabric Software/ Verbs	Verbs fails with segmentation fault error.	10.0.1
132062	HFI Host Fabric Software/MPI	When running mpis_stress test with MVAPICH, MVAPICH may fail with "Internal Error" after several hours of execution.	10.0.1
132436	HFI Host Driver	Unable to enable pcie error reporting: -5	10.0.1
132506	Management tools	Customers with unique sub-vendor and sub-device IDs experience a failure of the hfi_pkt function in the hostverify.sh script.	10.0.1
132525	Fabric Management/ PM	If you are running in a host subnet manager-embedded subnet manager configuration, a Virtual Fabric cannot run with a MaxMTU value of either 8192 or 10240. This causes checksum errors and forces one of the subnet managers into an inactive state, even though the XML configurations appear to be valid.	10.0.1
132584	Fabric Management/ FastFabric	FastFabric utility opaswitchadmin shows the status of power supply one (PS1) of unmanaged switch as "N/A" even if it is present and plugged in. It should report "ONLINE" when PS1 is plugged in and "OFFLINE" when it is not.	10.0.1
132709	Switch Fabric Management/SM	When running with multiple FMs, the standby FM's PM will have high CPU utilization if running with the PM on, and default settings for ImageUpdateInterval and PM SweepInterval. This behavior is magnified if the standby SM is an ESM, where the console can appear locked up because of the high CPU utilization.	10.0.1

Table 2-1. Resolved Issues (Sheet 2 of 2)

ID	Component	Description	Resolved in Release
132800	Switch Software/ Chassis Viewer	In Chassis Viewer, FRU information window does not display correct Manufacturer Date and Time information.	10.0.1
132903	HFI Host Fabric Software/HFI Host Driver	Packets sit in the highest priority queue for much longer than expected. Error message "user_mad: duplicate packet: both not response" is returned.	10.0.1
132937	HFI/HFI Pre-boot software	BIOS with support for UEFI Firmware Management may hang during boot if version 0x14 of the HFI UEFI Option ROM is installed on the HFI card.	10.0.1

2.3 Open Issues

Table 2-2 lists the open issues for this release.

Table 2-2. Open Issues (Sheet 1 of 4)

ID	Component	Description	Workaround
127520	Host Fabric Software/ MPI/PSM2	A segmentation fault may occur when running MVAPICH2 benchmark jobs utilizing greater than 72 mpi ranks per node.	Set MV2_USE_SHARED_MEM=0 in the MVAPICH2 command line.
129563	HFI Host Fabric Software/ MPI	Memory allocation errors with Mvapich2-2.1/Verbs.	When running mvapich2 jobs with a large number of ranks (for example, between 36 and 72 ranks), you must set the following parameters in /etc/security/limits.conf: * hard memlock unlimited * soft memlock unlimited Also, you must increase the lkey_table_size:LKEY table size in bits (2^n, $1 <= n <= 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to Appendix A in the Intel® Omni-Path Fabric Performance Tuning User Guide.
130233	HFI Host Pre-boot Software	Intel® Omni-Path DHCP fails to ID system based on hardware ID.	Workaround not available.
130435	HFI Host Driver	Packet traffic stops after reconnecting a Fabric Manager HFI port to a management allowed port that was previously marked "management not allowed".	Reboot the Fabric Manager node.
130791	Host Fabric Software	Using opapacketcapture with large amounts of traffic can cause performance issues on the given host.	Intel recommends you use opapacketcapture on hosts with lower packet rates and bandwidth.
131058	HFI Host Fabric Software	Running a job with Intel® MPI v5.0.2.044 using the OFA fabric (indicated by -IB or I_MPI_FABRICS=ofa) fails to start any test when the number of process ranks requested is greater than 36. The following error is generated: [10] Abort: Got completion with error 12, vendor code=0, dest rank= at line 858 in file ../../src/mpid/ch3/channels/nemesis/netmod /ofa/ofa_poll.c	Use one of the following workarounds: • Use Intel® MPI version 5.1.2, which is supported by this OPA Software Release. • Use the TMI fabric to invoke the higher-performing and more reliable PSM2 interface to the OPA hardware.

Table 2-2. Open Issues (Sheet 2 of 4)

ID	Component	Description	Workaround
131745	HFI Host Fabric Software/ MPI	When running OpenMPI 1.10.0 on SLES* 12 with large number of ranks per node (over 40), it may happen that the ORTE daemon (orted) "hangs" during the finalization of job.	Stopping and resuming the "hung" orted process allows the job to finish normally. To find the hung process, run the ps and find a node with several job zombie processes. In that same node, identify the orted process ID and send a stop signal (kill -19 <PID>) and a continue signal (kill -18 <PID>).
131934	Fabric Management Software	Redundant FMs may fail to synchronize and may log dbsync errors related to PM history redundancy.	Use the same value for the PM ShortTermHistory.StorageLocation setting in opafm.xml on all redundant FMs.
131947	HFI Host Driver	There is no support for HFI1 counters via: /sys/class/infiniband/hfi1_0/ports/1/counters	Use one of the following methods: • Use opainfo to get basic port information. • From a management node, use Intel® FastFabric commands. • Use the Intel® Fabric Manager GUI.
132127	HFI Host Driver	Certain server BIOS does not recognize HFI and fails to appear as a "UEFI Network" in the standard UEFI boot option.	Contact Intel Customer Support for details.
132197	HFI Host Fabric Software/mvapich2/ Verbs	The following mpi apps hang when run on mvapich2/verbs: • mpi_stress • osu_alltoall • osu_multi_lat The app osu_mbw_mr does not hang, but performance decreases as message size increases.	Set piothreshold=0 For instructions on setting module parameters, refer to Appendix A in the Intel® Omni-Path Fabric Performance Tuning User Guide.
132207	Open Software	Kernel crash caused by the ib_srpt module.	Install this kernel patch: https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=51093254bf879bc9ce96590400a87897c7498463
132345	Host Fabric Software/ MPI	When using Open MPI with the SLURM PMI launcher and PSM2, you may see the following error: Error obtaining unique transport key from ORTE (orte_precondition_transports not present in the environment). This is expected behavior. For details, see the Intel® PSM2 Programmer's Guide, in the description of psm2_ep_open() function, unique_job_key parameter.	To resolve this error, set the Open MPI environment variable OMPI_MCA_orte_precondition_transports in every node running the job. The format is 16 digit hexadecimal characters separated by a dash. For example, OMPI_MCA_orte_precondition_transports= 13241234acffedeb-abcdefabcdef1233 This key is used by the PSM2 library to uniquely identify each different job end point used on the fabric. If two MPI jobs are running on the same node sharing the same HFI and using PSM2, each one should have a different key.
132617	Fabric Management Software/Fabric Manager	Fabric Manager SwitchCascadeActivateEnable configuration setting can leave IPoIB in a non-working state for HFIs. This problem occurs when the Sm.SwitchCascadeActivateEnable setting in opafm.xml is enabled for all devices, and the FM is restarted. The restarted SM never issues a client reregistration of SA subscriptions request to endnodes.	Change the default Sm.SwitchCascadeActivateEnable setting in opafm.xml to a value of 0.

Table 2-2. Open Issues (Sheet 3 of 4)

ID	Component	Description	Workaround
132623	Fabric Management/ PM	Due to how the Performance Agent (PA) builds PA responses, a port may exist in one PA query and not in another. This can depend upon several factors, for example: • PmPortGroup • Internal/External membership • A possible Performance Manager (PM) Query Failure and/or a possible PM topology failure)	Workaround not available. ***Note:*** During a sweep if the failed ports count is a non-zero value, some data may be incomplete or missing.
132643	Host Software/Tools	hfi1_control -v Parameter does not work.	Workaround not available. Parameter to be removed as it is no longer needed.
132662	HFI Host Driver	During EFI boot, an EFI driver on the adapter is launched that returns error messages related to invalid temperature sensor readings.	The error messages can be ignored because the temperature sensor is being queried before it is ready.
132718	Software Installation	The Intel® Omni-Path SLES 12.1 package contains infiniband-diags version 1.6.4 which does not include the rdma-ndd service.	Install infiniband-diags version 1.6.5.
132760	DC Link Software	Link width downgrade to 0x width on TX [tx 0x0 (0), rx 0xf (4)] (AOC Link goes down.)	Contact Intel Customer Support for details.
132812	Switch Fabric Management	A GUI-less Linux* install (default option) does not install the bc package which is needed for some Fabric Manager scripts. The issue manifests itself by failing the hfi_pkt test and the log shows: bc command was not found.	Install the bc package.
133035	HFI Host Pre-boot Software	HFI UEFI boot asserts.	Do not run HPE BIOS with debug ASSERTs explicitly enabled.
133038	HFI Host Driver	PSM fails to acquire a context despite the kernel having contexts available. The following error message is returned: assign_context command failed: Device or resource busy	Reload hfi1 driver or reboot node.
133080	HFI Host Driver	Serial number reported using the hfi1_control -iv command is different than expected.	Workaround not available. Be aware that the serial number reported by hfi1_control is different from the serial number on the PCB assembly.
133114	HFI Host Driver	In some cases while the link is down, the HFI port may report the following in /var/log/messages: 'handle_qsfp_error_conditions: Cable RX channel 1/2 power too low' (AOC link is down.)	This message is benign and can be ignored.
133133	HFI Host Pre-boot Software	PXE boot timeout results in system hang.	Workaround not available.
133145	HFI Host Driver	hfi1_eprom has options to write, read, and erase information from fabric EPROM. However, it does not have an option to display current versions of the OptionROM version, Bootcode version, and Configuration file version. It is important for users to have this option to verify successful update.	Refer to Technical Advisory TA0008 for details (IBL #566152).
133238	Host Software/Verbs	Kernel panic in RHEL 7.2 running IPoSTL traffic.	Set piothreshold=0 For instructions, refer to Appendix A in the Intel® Omni-Path Fabric Performance Tuning User Guide.

Table 2-2. Open Issues (Sheet 4 of 4)

ID	Component	Description	Workaround
133241	HFI Host Driver	Under some circumstances, an HFI adapter returns QSFP memory read errors and is unable to bring up the physical link. (AOC link does not come up.)	Unplug and replug the QSFP module. If failure persists, recovery requires node reboot. Some cases may also require removal of the AUX power from management circuit on HFI adapter.
133269	Host Software/Verbs	GFP_NOIO error message using GFP_KERNEL when using IPoB on OPA in connected mode.	Workaround not available.
133371	Management Tools/FastFabric	In multi-rail fabrics, opaenableports and opadisableports do not work as expected when supplied with multiple egress ports. The scripts process the entire input file against the first egress port specified, leaving no data to be applied against other egress ports.	Call the script once for each egress port. For example, instead of this: opadisableports -p '1:1 1:2' 'dead servers' < disable.csv Issue two separate commands (one for each egress port) as follows: opadisableports -p '1:1' 'dead servers' < disable.csv opadisableports -p '1:2' dead servers' < disable.csv
133377	HFI Host Driver	irqbalance settings are not being honored correctly after a reboot.	Restart irqbalance after the hfi1 driver has loaded.
133408	HFI Hardware/PCIe card	In rare scenarios, loading the hfi1 host driver causes syslog messages including text such as "kernel:Uhhuh. NMI received for unknown reason 39 on CPU 0." On some systems, this may also be reported as "PCIe Fat" or as "CPU ERR2".	The error reporting is a transitory condition that arises on driver load. You can ignore the error and continue. Alternatively, ensure that the UEFI driver is being executed. This will perform the initialization before the OS boots and will avoid this error reporting. (For details, see the Intel® Omni-Path Fabric Host Software User Guide.)
133676	Fabric Management Software/Fabric Manager	Creating a virtual fabric with standby enabled causes opapaqueries to fail.	Do not enable standby on virtual fabrics.

§ §